The Effects of ShotSpotter Gunshot-Detection Technology on Shooting Fatality Rates

Terence Chau[*] Philip J. Cook[†] Henry Josephson[*]
Javier Lopez[*] Jens Ludwig[‡] Alejandro Roemer[§]
Adam Soliman[¶]

June 4, 2025

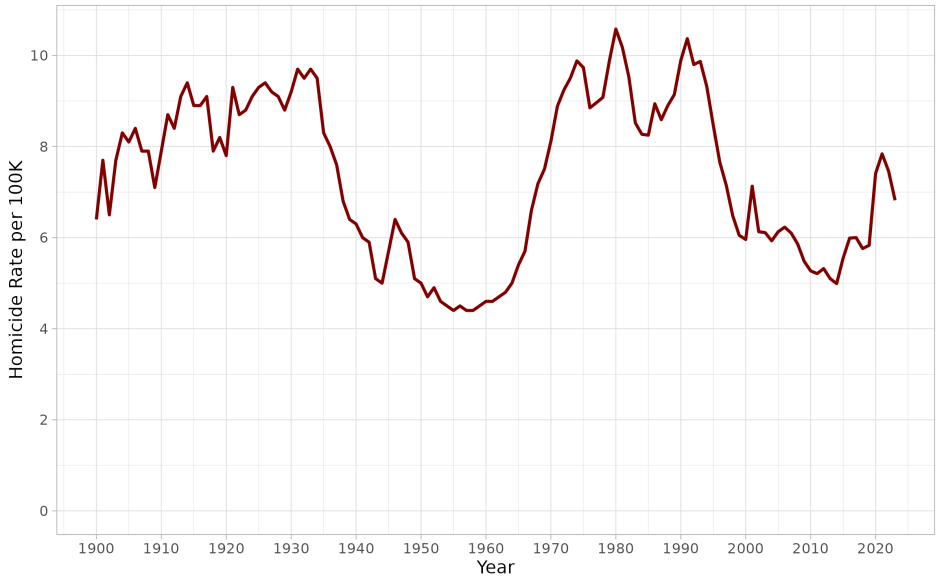
This project was supported by University of Chicago Crime Lab general operating funds. We have no financial or other conflicts of interest to report. Any errors and all opinions are our own. Contact: Jens Ludwig (jludwig@uchicago.edu)

[*]University of Chicago Crime Lab
[†]Duke University and NBER
[‡]University of Chicago and NBER
[§]Oxford University
[¶]Clemson University

I Introduction

Gun violence is a uniquely American problem among the rich countries of the world (Zimring and Hawkins, 1997). The public health impacts are sizable and enormously regressive. The problem has also proven to be hard to solve; while most mortality rates from almost every leading cause of death have plummeted since 1900, the rate of homicides in America—the overwhelming majority of which are committed with firearms—is almost exactly the same today as it was 125 years ago (Figure 1). This persistence comes in part from the limits of conventional wisdom on this problem and the resulting scarcity of effective, scalable solutions (Wellford, 2005; Cunningham et al., 2023; Ludwig, 2025).

Figure 1: U.S. Homicide Rate per 100,000, 1990 - 2023



Note: Data are sourced from the CDC's National Center for Health Statistics and compiled by the University of Chicago Crime Lab. Homicides from 1900–1948 include legal and police interventions, which constitute only a small proportion of all homicides. Rates from 1900–1932 are adjusted for registration area underreporting using the rates from Eckberg (1995).

The present study considers a technological intervention: ShotSpotter, a directional acoustic microphone system designed to detect gunfire and alert law enforcement. Although the technology has proven to be highly scalable, having been adopted now by 150 cities around the world (Topper and Ferrazares, 2024), the effectiveness of the tool remains the subject of considerable debate.

Previous research has mainly focused on the potential effects of ShotSpotter on law enforcement outcomes. The intended benefit of ShotSpotter comes from the fact that citizens may not call 911 in response to every instance of gunfire, or may report with either a delay or an error in identifying the gunshot's origin. By potentially alerting police more quickly and with greater accuracy, ShotSpotter creates a pathway to increase the fraction of shootings that result in an arrest (the "clearance rate"), or potentially,

an increase in gun confiscations. However, empirical evidence to date finds that, beyond improvements in response times, the impacts of ShotSpotter on criminal justice outcomes are limited at best (see, e.g., Piza et al., 2024; Cook and Soliman, 2024; Topper and Ferrazares, 2024).

The present study considers ShotSpotter's potential effects on a different outcome of particular importance for public health: shooting fatality rates. The motivation for this analysis is the possibility that quicker police responses to shootings could potentially translate into more timely medical assistance to victims. We use data from Chicago, where at its peak the technology was in operation in a set of police districts that covered around 50% of the city's land area that includes around 83.7% of all shootings. We use a regression discontinuity (RD) design that exploits two sources of variation: 1) *Time*, specifically, the phased rollout of ShotSpotter across various police districts starting in 2017, followed by the abrupt, citywide termination of the system on September 23, 2024 when Chicago ended its ShotSpotter contract, and 2) *Geography*, by comparing shooting fatality rates across the boundary of districts with and without ShotSpotter while the technology was in operation.

The present draft presents preliminary results that capitalize on the geographic variation only. The first set of results we present in this draft establishes the validity of the geographic RD variation that we use for identification. This design essentially compares shooting fatality rates "across the street"—that is, on one side of a police district boundary relative to the other. One might naturally worry that other determinants of shooting outcomes might vary discontinuously at these boundaries. For example, Chicago is famously a "city of neighborhoods" (Sampson, 2012) and thus social conditions or social norms may vary discontinuously across areas. Similarly, we might worry police districts themselves might differ in systematic ways that affect shooting outcomes. Yet in practice, we find that when we estimate the geographic RD design during the pre-ShotSpotter period, we see point estimates for impacts on outdoor shootings, the focus of our analysis, that are quite close to zero and consistently statistically insignificant.

In contrast, during the period after ShotSpotter was implemented in all treated districts, our point estimate of the number of gunshot victims who die suggests a discontinuous downward drop in shooting fatality rates on the order of 2.5 percentage points (equal to reducing the rate by 14.5%). While this research design seems to have good internal validity, the challenge is statistical power: Even with a proportionally meaningful impact estimate like this, the p-value is 0.286.

What should we make of these preliminary results in the meantime? Imbens (2021) argues that for policy decisions the standard p-value cutoff of 0.05 used for statistical testing of scientific hypotheses is not appropriate. Science is inherently conservative, so this stringent cutoff reflects understandable caution in overturning conventional wisdom. But that same approach for policy would create a strong bias toward the status quo. Instead, Imbens argues, policymakers should focus on something closer to the expected value of the candidate policy by considering the point estimate, its uncertainty, and the implied magnitudes of the benefits and costs.

If we were to apply that framework to our setting, the most obvious cost would be the annual fees to the city of Chicago over the life of the ShotSpotter contract, which have averaged between $7 and $10 million per year (in 2024 dollars). There is also a reduction in arrests for non-gun crimes as documented by Topper and Ferrazares (2024). Two additional unintended consequences include the possibility of overly aggressive police responses to false positive ShotSpotter alerts, i.e., an alert where there was no weapon discharge, and a concern by some city residents over the extra electronic surveillance that ShotSpotter creates in Chicago's most disadvantaged and segregated areas where gun violence is disproportionately concentrated. We unfortunately have no way to quantify the magnitude of those potential costs.

On the potential benefit side, there were 1,966 outdoor shootings in Chicago in calendar year 2024 in those police districts that had ShotSpotter active. A 2.5 percentage point impact on shooting fatality rates would imply 49 fewer fatalities. If the city of Chicago used the same value per statistical life the federal government uses for safety regulations—$13 million—then our estimates imply that for $7-10 million per year (plus the harder-to-quantify additional costs described above), our point estimate taken at face value would imply the social value of reduced fatality rates would be on the order of $637 million.

Whether the potential for that benefit, given the costs and degree of statistical uncertainty with our preliminary impact estimates, is worth the investment is ultimately a normative question for policymakers to decide. Our hope is that incorporating additional data analysis, as we have done here, may help narrow down the uncertainty about the technology's tradeoff of benefits and costs.

The remainder of the paper is structured as follows. Section 2 reviews the existing literature, while Section 3 explores potential mechanisms. Section 4 describes our data and the Chicago context in which the evaluation takes place. Section 5 outlines the empirical methods, and Section 6 presents the results. Section 7 discusses the policy implications of these findings.

II Previous research

Most of the existing social science research on ShotSpotter's effects has focused on criminal justice outcomes like arrests, guns confiscated, or police response times.

Cook and Soliman (2024) study the effects of a one-year ShotSpotter pilot deployment over three square miles of Durham, North Carolina from December 15, 2022 through December 14, 2023.[1] The police department adopted a protocol in which ShotSpotter

[1]Over this period the police department received 1,447 total notifications of gunfire in the pilot area, of which 57% were ShotSpotter alerts with no 911 call, 28% were 911 calls with no ShotSpotter alert, and 15% had both 911 and ShotSpotter alerts (so the technology doubled the number of gunshot notifications).

alerts would be treated as Priority 2 incidents—that is, immediate deployment but without sirens and flashers. Response times to shooting incidents were measurably reduced compared with what they would have been if the police had had to rely on a 911 call. And in a number of validated shooting cases, there was no 911 call at all. Indeed, the police made 7 shooting-related arrests where law enforcement learned about the call only from ShotSpotter, i.e., with no corresponding 911 call. Whether that is a net increase of 7 arrests relative to the counterfactual scenario of no ShotSpotter depends on whether the police would have found out about those shootings in the future through some other means. The cost of the technology in the pilot area over this period was about $200,000.

Topper and Ferrazares (2024) study the staggered roll-out of ShotSpotter across police districts in Chicago starting in 2017. They find that holding total police manpower fixed, adoption of ShotSpotter slows down dispatch times to non-ShotSpotter-related calls by 22% and reduces the probability of an arrest in those calls by 9%. In general, the empirical evidence finds that ShotSpotter does not reduce aggregate gunshot-related victimizations in Chicago, nor does it increase gun-related arrests (Piza et al. (2024), Topper and Ferrazares (2024)).

Connealy, N.T., Piza, E.L., Arietti, R.A. et al (2024) also study the effects of ShotSpotter adoption in Chicago, using a synthetic control design. They find the technology increased the number of gun confiscations by 12%. There is no detectable impact on other criminal justice outcomes examined, although many of the estimates are not statistically precise. Connealy, N.T., Piza, E.L., Arietti, R.A. et al (2024) also look at fatal and non-fatal shootings and find no detectable effects.

There have been evaluations of ShotSpotter on police activity and violence in other cities with mixed results, such as in Houston (Cheng, 2023), Winston-Salem (Mares and Buettner, 2024), Kansas City (Piza et al., 2024), St. Louis (Mares and Blackburn, 2021), Camden (Goldenberg et al., 2019), and Las Vegas (Koren, 2018). These studies suggest that while ShotSpotter may improve certain aspects of police operations, particularly response times and evidence collection, its direct impact on reducing gun violence remains uncertain. In a similar vein, (Doucette et al., 2021) analyze homicide rates in large metropolitan counties from 1999 to 2016, exploiting the adoption of ShotSpotter in some areas during this period, but find the estimated effects on homicide rates are again imprecise.

The research literature to date, in other words, has focused far more on enforcement outcomes than on the public health outcome of shooting fatality rates (defined as the ratio of fatal shootings to total gunshot injuries). Identifying that impact is the goal of the present paper.

III Mechanism and Logic

Our focus in this analysis is on the effect of the ShotSpotter installation on a single outcome: the likelihood that a gunshot victim survived. In recent years in Chicago, there has been about one death for every 6 victims of assault who suffer a gunshot injury (Ludwig and Miller, 2024). The likelihood of death is influenced by a number of factors, such as whether the victim is shot in a particularly vulnerable body part, such as the torso or head, rather than an extremity, the number of gunshot wounds they suffer, and the caliber of the bullet(s) (Braga and Cook, 2018).

In cases where the victim does not die immediately as a result of their injuries, the likelihood of death also increases with the amount of time before first aid is administered. One correlational study found that the fatality rate for gunshot wounds was about 2% higher for every additional minute until emergency medical services arrived on scene (Nasser et al., 2020). In many cities, including Chicago, the police are authorized and trained in administering first aid, are equipped with first aid kits (including tourniquets, chest seals, pressure bandages, and hemostatic gauze), and are often the first responders to gunshot victims.[2] While police officers do not transport victims to the hospital, they can mobilize emergency medical response.

The evidence is clear that ShotSpotter reduces police response time in shootings. In some cases the ShotSpotter alert is the only notification that the CPD receives. Even for incidents in which there is a 911 call, it is typically true that the ShotSpotter alert reduces response time: these alerts tend to be received more quickly than 911 calls (Cook and Soliman, 2024), and they typically provide more precise information about the location of the shooting, thus reducing search time by the responding officer. A quicker response to the victim provides a greater chance for officers to stop the bleeding and summon an ambulance.

Of course there are many cases in which the victim survives without any medical assistance in the field, and others where the victim dies immediately. It is the intermediate cases—serious injuries in which the victim survives but would die within minutes in the absence of first aid—in which there is a possibility that the police response time matters.

This simple conceptual framework highlights the limitations of most existing discussions of Shotstpotter's effects on shooting fatality rates, which have focused on descriptive tabulations of what share of shootings involve a ShotSpotter alert, a 911 call, both or neither.[3] However, this type of analysis is not very informative about public health impacts because not every shooting where there is only a ShotSpotter alert leads to a saved life, and because ShotSpotter could in principle help save lives even in cases where there is a 911 call, if the ShotSpotter alert gets police to the scene faster than

[2]Chicago's Civilian Office of Police Accountability noted in 2018 that this sort of training and equipment is considered best practice nationally; see Banks (2018). The department has since taken steps in this direction, see for example Chicago Police Department (2023).

[3]See for example, Ferguson and Witzburg (2021).

they would have through quicker notification or more accurate location information.

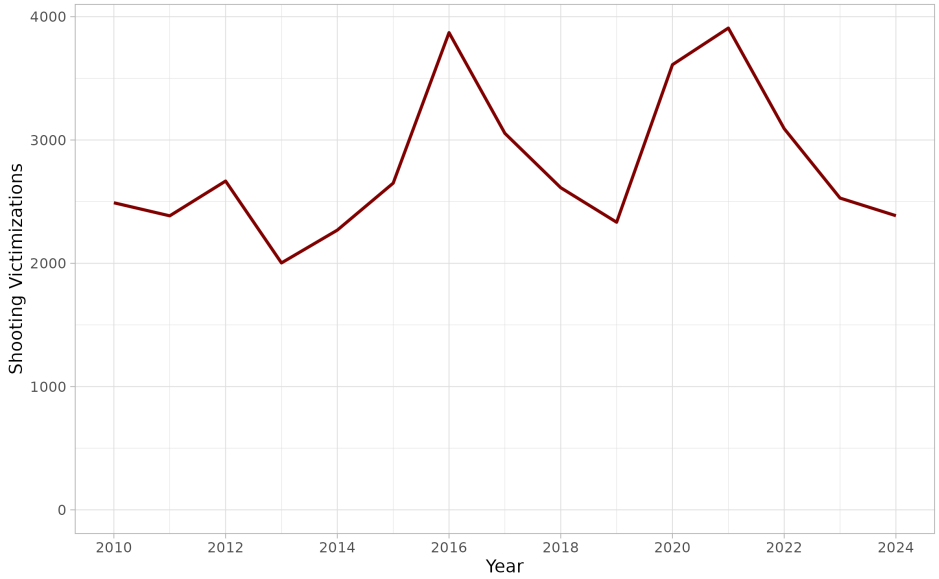
In other words, to understand the technology's impacts on shooting fatality rates a different research design will be required. That is the contribution of the present paper.

IV Data and Context

We rely on the public victimization data from the Violence Reduction Dashboard (VRD), hosted by the Chicago Mayor's Office and created with technical assistance by our research center, the University of Chicago Crime Lab.[4] The VRD provides up-to-date information on violence and victimization incidents at the victim by incident level, and it updates daily with an approximately 48 hour lag. Data fields in the VRD include information such as whether the incident was fatal (the main outcome of interest), the type of incident, the latitude and longitude coordinates, and the location description of the incident which allows us to distinguish if the incident occurred indoors or outdoors. The data were last queried on January 27, 2025.

In total, the VRD data contains 41,347 outdoor shootings at the victim-by-incident level in Chicago from January 1, 2010 through September 22, 2024, with 22,927 taking place after the first ShotSpotter implementation.[5]

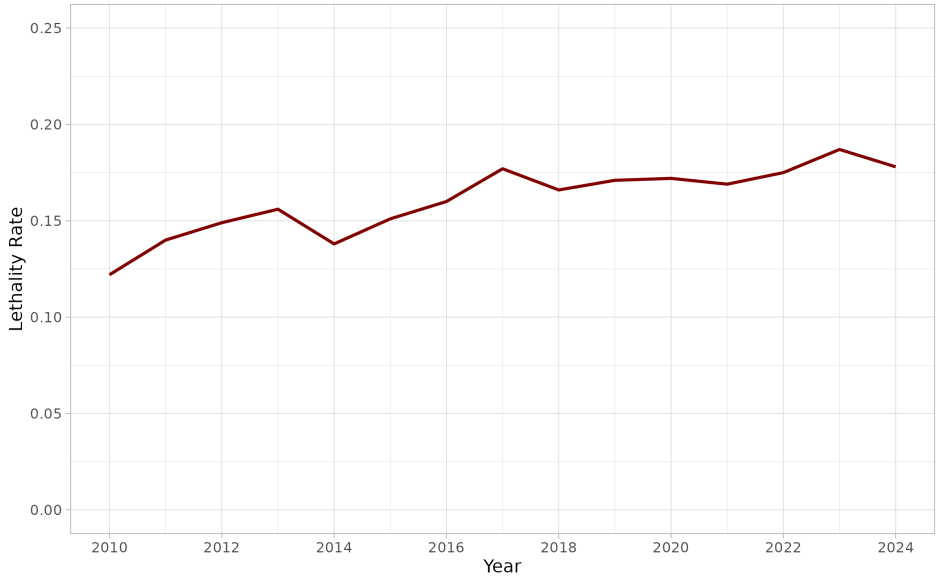
Figure 2: Yearly Gunshot Injury Trends in Chicago, 2010-2024



[4]The underlying data is hosted at the Chicago Data Portal and the VRD can be used to visualize it.

[5]Outdoor shootings as well as those that may have occurred outdoors but are not clearly encoded (i.e. gas stations, parking lots, Chicago Transit Authority platforms etc.) are the focus of the main analysis sample and summary statistics, unless noted otherwise, as they consist of 88.8% of all shootings in our dataset and most plausible to be detected by ShotSpotter.

Figure 3: Yearly Fatality Trends in Chicago, 2010-2024



The 2010-2024 period observed 41,862 outdoor shooting victimizations as seen in Figure 2. Outdoor shooting victimizations have increased in the sample period from a low of 2,003 in 2013 to a high of 3,908 in 2021. The overwhelming majority of gunshot victims are male (87.3%). Victims were also disproportionately likely to be Black non-Hispanic (77.2%), and young. While a disproportionate share of public attention goes to adolescents, who are 23.2% of victims, a very large share of victims (43.4%) are in their 20s and many even in their 30s.

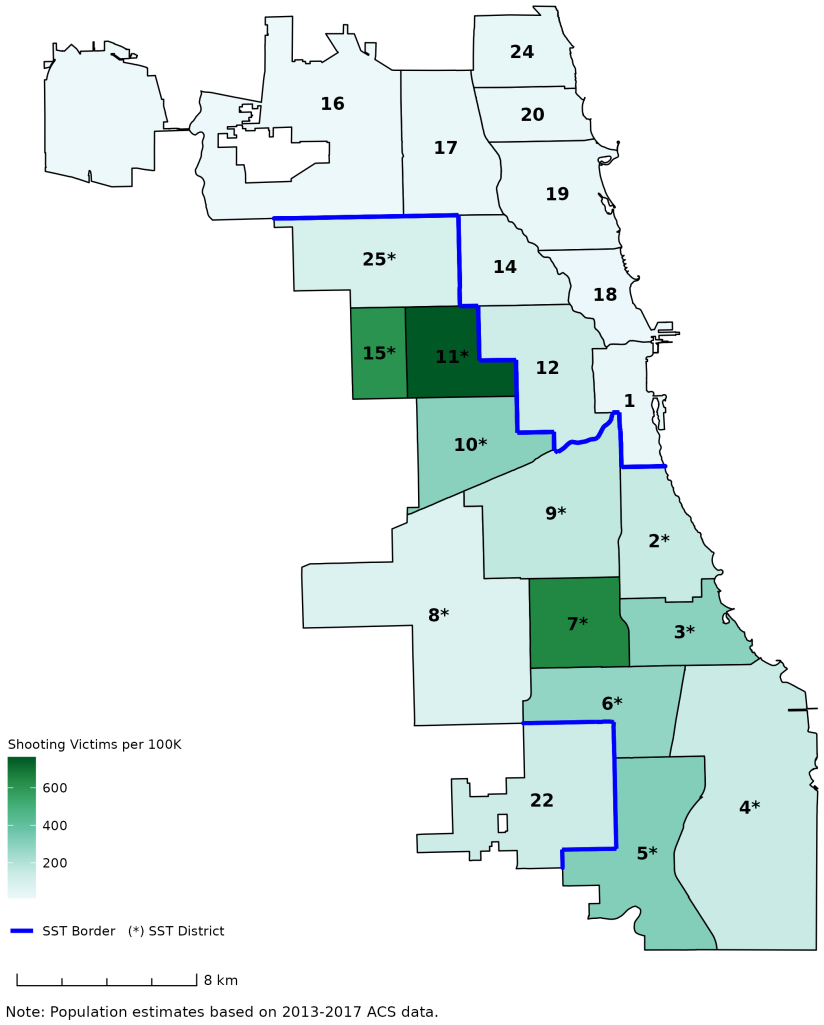
Gun violence has been getting deadlier in Chicago over time. Over the 2010-2024 period the average fatality rate is 16.1% of outdoor shootings, but that figure has increased over this 15 year period by fully 5.6 percentage points (Figure 3). This increase is analogous to approximately 134 additional lives lost per year relative to what would have happened under 2010's fatality rate. One hypothesis for why the shooting fatality rate has been increasing that enjoys some support in the data is a growing use of Glock switches and bump stocks, which modify semi-automatic firearms to fire automatically (like machine guns), combined with growing use of high-capacity magazines that determine the number of rounds that can be fired before reloading (Ludwig and Miller, 2024).

V Methods

The city of Chicago began adopting ShotSpotter in 2017 on the heels of a large surge in gun violence in 2016 as shown in Figure 2. The technology was prioritized for police districts on the city's South and West Sides where gun violence is concentrated (see, for example, 2016 outdoor victimization rates prior to ShotSpotter implementation in Figure 4). Then at midnight of September 22, 2024, the city of Chicago allowed its contract with ShotSpotter to terminate and the technology went offline citywide all at

once, with the last alert occurring at 11:54pm the night of September 22. These are the different sources of variation that we rely on for identification.

Figure 4: Shooting Victimizations per 100K by CPD District, 2016



Note: Population estimates based on 2013-2017 ACS data.

A naive comparison of shooting fatality rates would simply compare outcomes in those districts with and without ShotSpotter during the treatment window, ignoring the fact that the ShotSpotter areas are the most violent ones in the city with a high volume of shootings (Figure 4). All else equal, the higher demand for medical personnel and police can potentially lead to tighter police resource constraints and hence higher response times on average, which in turn could lead to higher fatality rates. This means that if we run a naive regression comparing the ShotSpotter to non-ShotSpotter districts overall as in Table 1, we see fatality rates that are higher in the ShotSpotter areas, but this cannot be reliably interpreted as evidence about the causal effects of ShotSpotter per se.[6]

[6]Similarly, a simple comparison of fatality rates in shooting incidents that did versus did not trigger ShotSpotter can also be confounded by various factors. For example, many shootings that trigger a ShotSpotter alert will also independently produce a 911 call reporting the incident, e.g., if ShotSpotter sensor placements are installed in more population-dense areas. In the counterfactual absence of the

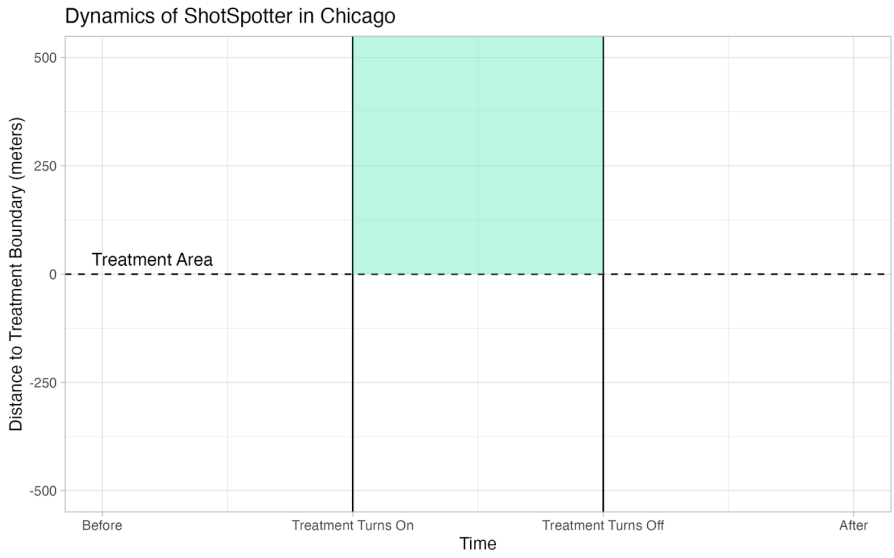
Table 1: Naive Comparison of ShotSpotter on Fatality Rates, All Districts

	Control Mean	Treatment-Control Diff.	p-value	N
Fatality Rate	0.160	0.017	0.021	19,198

Note: Robust standard errors. Each observation is a shooting victimization.

Figure 5 illustrates the different sources of identifying variation we use in our study. The y-axis in the figure represents geographic variation, specifically distance from the boundaries between ShotSpotter districts (top or >0) and non-ShotSpotter police districts (bottom or <0). One potential complication here is that sound travels; microphones in ShotSpotter districts sometimes could detect gunfire in nearby districts that lack the technology. But in practice ShotSpotter only relays gunfire notifications to the Chicago police department from shootings that the technology's directional acoustic microphones detect in the ShotSpotter districts themselves (with a very small buffer near the district boundaries, as we document below). Time is indicated on the x-axis; part of our identifying variation comes from the turning-on of ShotSpotter starting in 2017 (moving from the "before" period to intervention period across the first time threshold to the left), and then the turning off of the technology again in 2024 (crossing the second time threshold to the right).

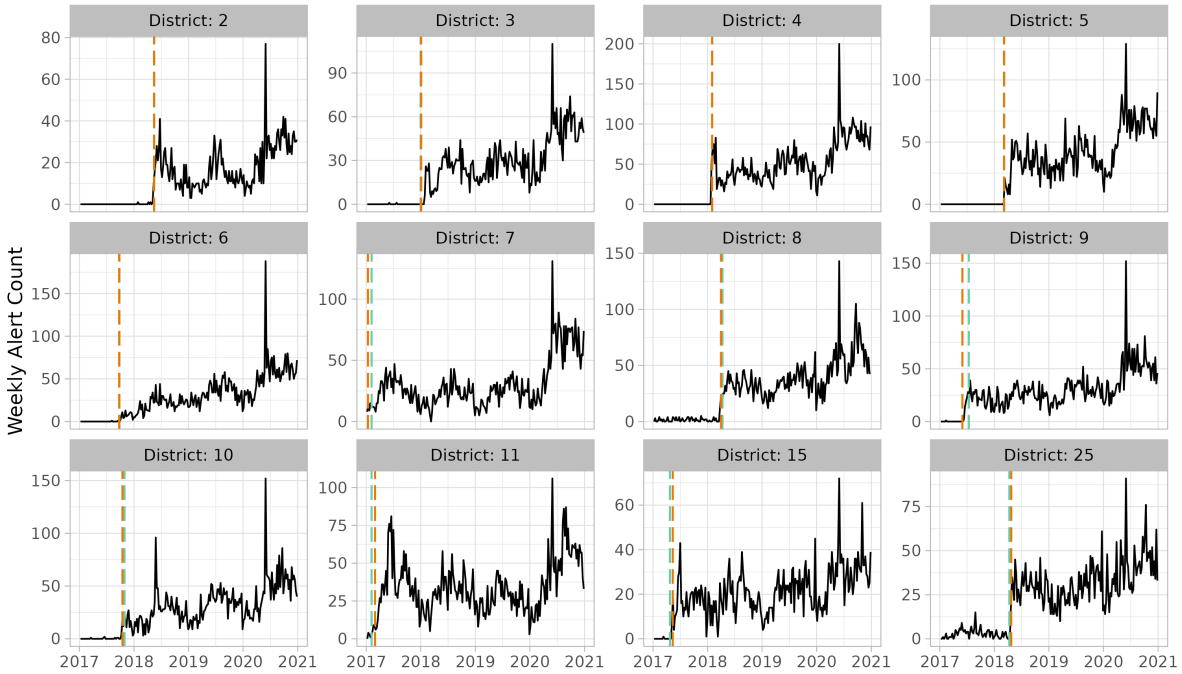
Figure 5: Two-dimensional regression discontinuity design using both time and geography



One potential complication with this source of identifying variation from the time dimension stems from conflicting dates in the ShotSpotter alert data and the research

ShotSpotter alert, a 911 call would have still been generated, reducing the lethality of the shooting regardless. This would lead to an overestimate of the effectiveness of ShotSpotter, as its treatment effect is bundled with the effect of the 911 call. This particular concern is mitigated to the degree that reporting rates (as measured in the share of shooting incidents tied to a given 911 call) are uniform across the sample. In future work, we aim to merge 911 call data into our sample to assess the degree of this issue.

Figure 6: Weekly ShotSpotter Alert Counts by District



Note: Rollout dates from Piza et al. (2024) in teal and from Topper and Ferrazares (2024) in orange.

literature about when the technology went live in different Chicago police districts; we set the rollout dates in our study in accordance with the installation dates in Topper and Ferrazares (2024), which we have verified are quite close to the dates in which CPD began to receive live ShotSpotter alerts in each district on a routine basis (see Figure 6 below).

A second potential complication with this source of over-time variation is that in some districts the technology went "live" to full implementation immediately; in other districts there was more gradual adoption.[7]

In our final analysis, we will use the approach from Cheng (2024) to generate our regression discontinuity (RD) estimate with both of these running variables—both geography (just on the other side of ShotSpotter and non-ShotSpotter district boundaries, which may be more similar areas than comparing the two types of districts overall) and time.[8]

In the current draft we use only the geographic variation; that is, we focus our analysis on the period in which ShotSpotter was in effect (the middle time band in Figure 5) and essentially compare shooting fatality rates across the border of ShotSpotter and non-ShotSpotter districts, building on the initial design from Black (1999). More

[7]As we intend to leverage the rollout and shutdown of ShotSpotter in future work, we will create estimates that use varying definitions of rollout to assess the robustness of our results. It is also the case that what counts as full ShotSpotter coverage within a district depends on the number of microphones installed in an area, population density, presence of tall buildings, and the topography of the terrain (for example, hilly versus flat).

[8]We expound on this design in Appendix Section A.V.

specifically, we estimate a standard sharp RD model that uses locally linear regression to model the relationship between the running variable and the outcome:

$$Y_i = \beta_1 + \beta_2 r_i + \beta_3 Z_i + \beta_4 f_h(r_i) + \epsilon_i \tag{1}$$

Where we define Y_i as an indicator for whether a shooting was fatal, r_i as the normalized distance to the police district boundary, Z_i as an indicator for whether the incident occurred in a district with or without ShotSpotter, and h as the bandwidth that defines the amount of data near the boundary that is used in the model. In this model, β_3 represents the intent to treat estimate at the boundary, $E(Y_{1i} - Y_{0i}|r_i = 0)$.

Following current best practices in the regression discontinuity literature, all regressions are estimated using local linear regression with a triangle kernel, implement data-driven bandwidth selection methods (Imbens and Kalyanaraman, 2012) to define the baseline bandwidth, and use inference methods robust to these bandwidth selection methods as outlined by Calonico et al. (2014) and Cattaneo et al. (2020). Estimates using a variety of bandwidths, kernels, polynomial specifications, and inference methods are presented as robustness checks (Figures 11 and A.3, and Appendix Table A1).

The estimate above is not necessarily the same as the effect of an incident having a ShotSpotter (SST) alert. This is because not all shooting incidents within an SST area receive an SST alert, and some incidents right on the control side of the boundary can potentially trigger an alert. In future work, we intend to analyze non-compliance and estimate a fuzzy regression discontinuity, where crossing the boundary serves as an excluded instrument for SST take-up.[9]

The main identifying assumption required for β_3 to be an unbiased estimator is that on average, potential outcomes in lethality with and without SST are continuous at the district boundary; that is, the lethality likelihood would not discretely jump at the police district boundary in the absence of ShotSpotter. This assumption need not hold if for example, police or medical personnel response times to shootings are discontinuously changing between neighboring police districts right at their border for reasons other than ShotSpotter (Topper and Ferrazares, 2024). Similarly, this identifying assumption could be violated in principle if officers in neighboring districts have different practices for rendering first aid to shooting victims; for example, even though all Chicago PD officers are supposed to be trained on first aid now and carry tourniquets and pressure bandages, one could imagine that in principle some districts are more supportive of releasing officers from duty to attend training, or more strict in ensuring officers have first aid kits. This assumption encapsulates the traditional continuity assumption in a regression discontinuity design. We also require the assumption that there is no manipulation in the density of shooting incidents near the geographic discontinuities. We discuss the evidence on these assumptions in what follows.

[9]Note: There are no unique identifiers to link shooting victimization records with SST alert data. As such, to establish take-up, we perform spatial-temporal linking by identifying the SST alerts closest to shooting victimizations within defined time and distance parameters. This process, along with an assessment of alternative matching parameters, is detailed in Appendix Section A.III.

Given the geographic RD design employed in this study, it is relevant to also discuss the accuracy of the latitude and longitude in the data. To maintain a victim's anonymity, points in the victimization data are randomly jittered by the City of Chicago by up to the size of an average city block. Specifically, a circle of this size is drawn around each exact victimization location in the data, and a random point is chosen within this circle as the reported latitude and longitude. Chicago is designed around a grid system. Due to this, the average block measures 1/8th of a mile (201 meters) on each side.[10]

The fact that the running variable is measured with random error should introduce both some bias and some variance into our estimates. The jittering of addresses should introduce some attenuation bias as some treated observations can be jittered into control areas and vice-versa. As a robustness check we leverage police district labels, which are not re-coded after jittering; those confirm that results are not driven by this random noise injected into the geocoding process—we find that only 250 (1.7%) of outdoor shooting victimizations within 2km of the boundaries of districts in our sample ranging from January 1, 2010 to September 23, 2024 are jittered from non-SST districts to SST districts or vice-versa. The variance or extra "noise" in our estimates comes from the fact that jittering of the shooting incident locations means that our attempt to control for the pre-existing continuous relationship between geographic distance and shooting fatality rates will have less explanatory power than it would absent that source of measurement error in r_i, which in turn should inflate our standard errors somewhat relative to a world without that measurement error.

VI Results

In this section we first show that there are no signs of a treatment effect across the geographic boundary during the *pre*-ShotSpotter period as a way to validate the identifying assumptions of our geographic RD research design. We then show that during the post-ShotSpotter period the results consistently suggest a negative estimate on shooting fatality rates (including at the optimal MSE bandwidth) although the data are somewhat noisy.

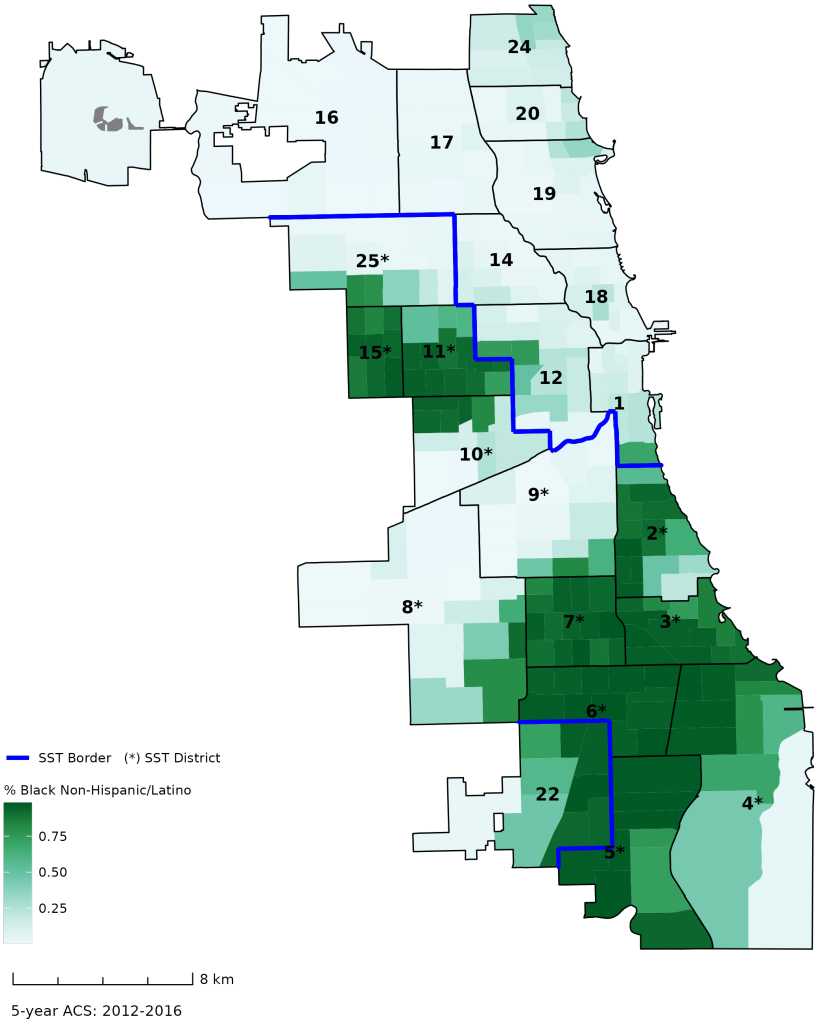
VI.1 Tests of Identifying Assumptions

We begin by examining the evidence for balance in pre-treatment covariates. Unfortunately, demographic data on the city of Chicago population as a whole, such as Census Block Group records, are not collected or reported in a way that precisely aligns with police district boundaries and any conversions would dilute Census estimates. Moreover, estimating imbalance at the bandwidth threshold using our main RD design would require tracking how socio-demographic characteristics vary within specific distances from

[10] See University of Chicago (2020) and Chicago Department of Transportation (2007).

district borders— a task that would necessitate geocoded *household*-level (as opposed to neighborhood-level) data, which the Census Bureau does not readily provide.

Figure 7: Estimated % of Black Residents by CPD Beat



However, we are able to carry out this type of continuity test for the characteristics of shooting victims specifically, as shown in Table 2. We draw on data during the pre-Shotspotter period, from January 1, 2010 to January 13, 2017. In general, we find that these characteristics are mostly continuous across the cutoff pre-treatment.[11]

A different test of the continuity assumption is to look at how gunshot fatality rates differ at the district boundaries *before* ShotSpotter goes into effect. We present results from estimating equation (1) using the same pre-treatment timespan of data, before ShotSpotter was rolled out in any district. There appears to be no pre-existing difference in fatality rates across study district boundaries (Figure 8 & Table 3). The point estimate is typically on the order of around 1-2 percentage points in the various reasonable choices of bandwidth and is not close to statistical significance under similar

[11]We replicated this exercise on the during-treatment window data (May 16, 2018 to September 23, 2024) and outline the results in appendix section A2.

Table 2: Continuity Tests of Victim/Incident Covariates, Pre-Treatment

Covariate	Control Mean	Treatment-Control Diff.	p-value	Observations
Age 0-19	0.259	-0.030	0.532	3511
Age 20-29	0.485	-0.009	0.799	3511
Age 30-39	0.176	0.047	0.171	3511
Age 40-49	0.057	-0.013	0.395	3511
Age 50-59	0.015	-0.002	0.735	3511
Age 60-69	0.007	0.006	0.215	3511
Age 70-79	0.001	0.002	0.184	3511
Age 80+	0.000	0.000	0.198	3511
Age (Missing)	0.022	-0.020	0.027	3599
API/Other NH	0.004	0.003	0.319	3484
Black NH	0.584	0.086	0.207	3484
Hispanic/Latino	0.374	-0.071	0.326	3484
White NH	0.037	-0.018	0.287	3484
Race (Missing)	0.029	-0.008	0.351	3599
Male	0.901	-0.000	0.884	3505
Sex (Missing)	0.023	-0.014	0.133	3599
Jan	0.059	-0.040	0.014	3599
Feb	0.031	-0.005	0.715	3599
Mar	0.076	-0.045	0.008	3599
Apr	0.082	-0.007	0.707	3599
May	0.088	0.085	0.000	3599
Jun	0.107	0.005	0.596	3599
Jul	0.113	-0.019	0.192	3599
Aug	0.114	0.047	0.002	3599
Sep	0.093	0.015	0.332	3599
Oct	0.093	0.001	0.775	3599
Nov	0.066	-0.016	0.216	3599
Dec	0.079	-0.020	0.411	3599
Monday - Friday	0.566	-0.025	0.220	3599
CPD Watch #1 11:00pm-7:00am	0.395	0.051	0.202	3599
CPD Watch #2 7:00am-3:00pm	0.142	0.013	0.845	3599
CPD Watch #3 3:00pm-11:00pm	0.463	-0.064	0.168	3599

Note: Each test carried out using the baseline specification outlined in equation (1), using a local linear regression, optimal bandwidths from the main RD, bandwidth selection robust inference, and a triangular kernel. Observation totals are comprised of treatment and control, within bandwidth.

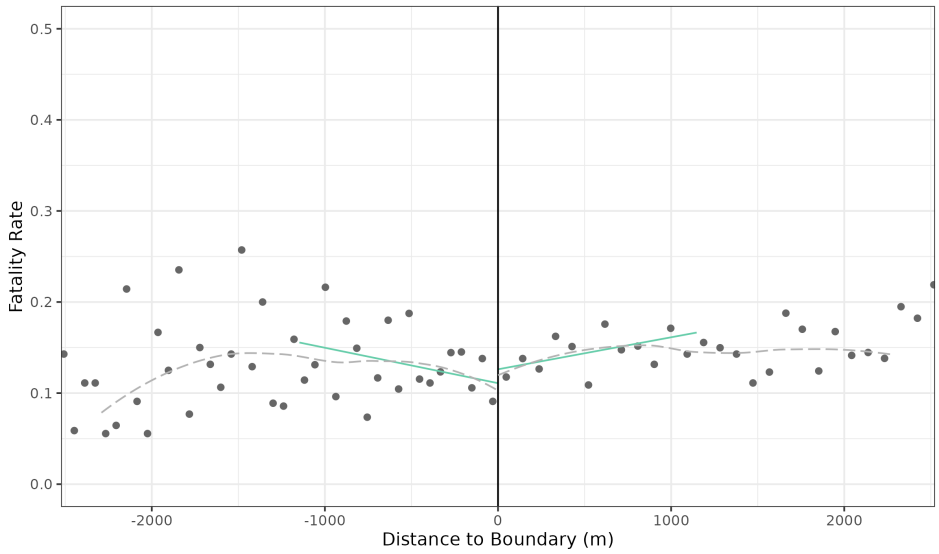
Table 3: Sharp Regression Discontinuity Estimates

	Outcome: Fatal Shooting	
	Treatment Period	Pre-Treatment Period
In ShotSpotter Area	-0.025 [-0.091, 0.027] (0.286)	0.015 [-0.039, 0.064] (0.634)
Control Mean in Bandwidth	0.17	0.128
Optimal Bandwidth	1032.105	1147.441
Observations	4021	3926

Note: Bandwidth selection robust confidence interval in brackets, p-value in parentheses. Estimates produced using local linear regression with a triangular kernel.

specifications as the basline RD ($p - value = 0.634$). This is the case across various bandwidths (Appendix Figure A.2).

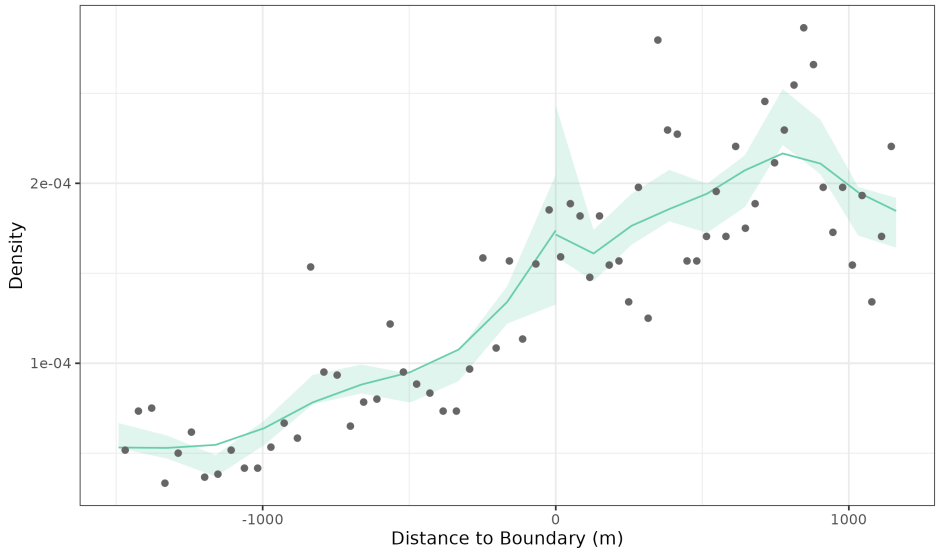
Figure 8: Placebo Regression Discontinuity Estimates in the *Pre*-ShotSpotter Period



Note: Optimal bandwidth local linear regressions in teal (coefficient = 0.015, robust p-value = 0.634), local linear smoother fitted on each side of the discontinuity in dashed gray. Bins created using mimicking variance evenly-spaced method (Calonico et al., 2015)

Of course one might wonder or worry whether there are other changes around the time that ShotSpotter is introduced that could affect some districts differently than others, and so potentially confound our RD estimates. A natural way to consider testing that possibility would be to generate post-ShotSpotter RD estimates for criminal justice outcomes that conceptually not be affected by the introduction of ShotSpotter itself. The problem with that sort of specification test is that there are few if any types of criminal justice outcomes that might not conceptually be affected by ShotSpotter's introduction given the spillover effects due to police capacity constraints, as outlined in Topper and Ferrazares (2024).

Figure 9: Running Variable Density Test



Note: p-value = 0.253. McCrary (2008) density test (shown in Appendix Figure A.1) p-value = 0.983. Data includes only observations within two kilometers of the district boundary.

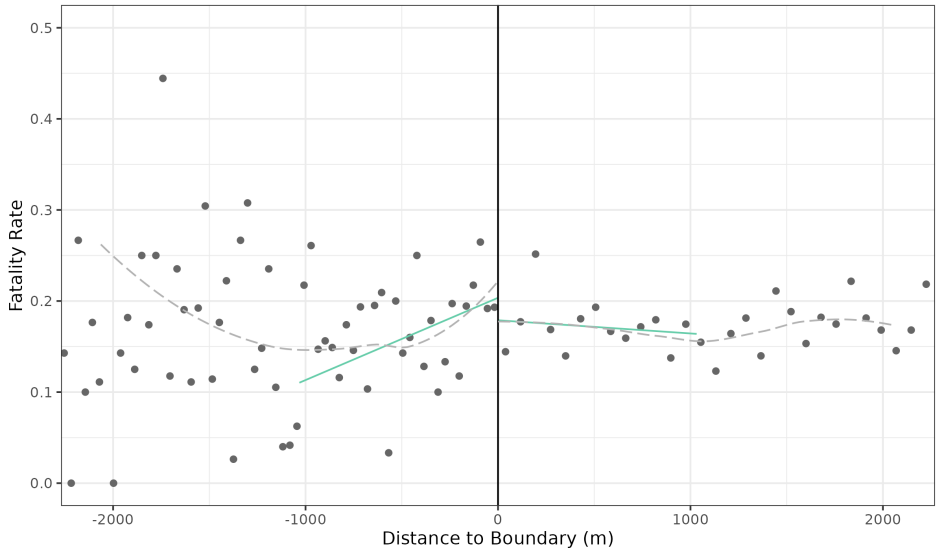
A final assumption needed for the regression discontinuity design to work is that there should be no manipulation in the density of shooting incidents across the cutoff during the treatment period. In Figure 9 we provide formal evidence of this via a Cattaneo et al. (2020) density test, along with McCrary (2008) density test in Appendix Figure A.1.

VI.2 Main Results

The results of estimating our main regression discontinuity design are plotted in Figure 10. Specifically, we plot evenly spaced bins of the running variable, distance to the district border, in the x-axis, against the fatality rate of shooting victimizations in each bin in the y-axis. Two sets of regression lines are plotted, a local linear regression within the estimated optimal bandwidth, and a local linear smoother within about 2km on each side of the border.

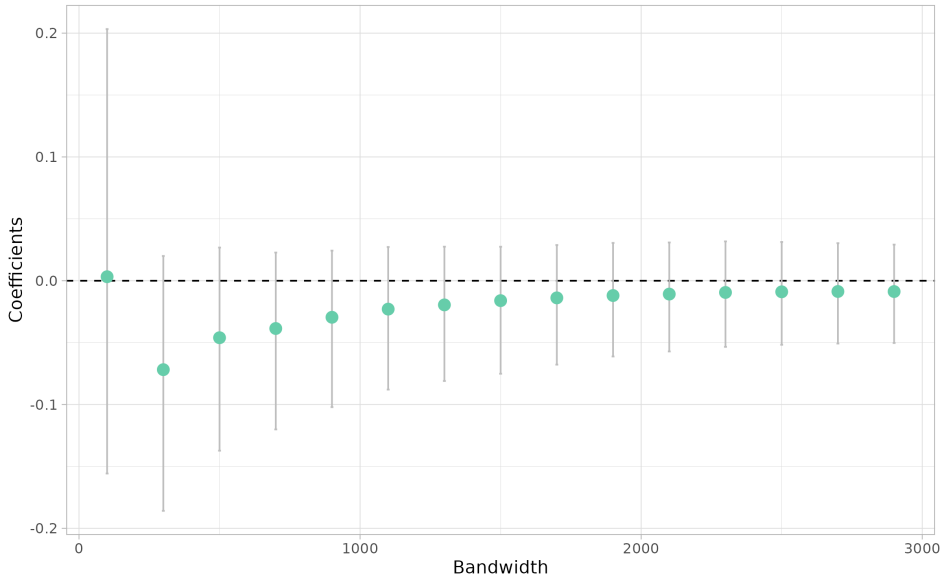
Table 3 shows regression estimates from the local linear specification, and Figure 11 shows equivalent estimates for all bandwidths from 100m to 2.9km from the border. Our baseline specification implies a point estimate of -2.5 percentage points with a 95% robust confidence interval ranging from (-0.091, 0.027). Generally for different specifications the point estimates are consistently negative but never statistically different from zero.

Figure 10: Sharp Regression Discontinuity Estimates



Optimal bandwidth local linear regressions in teal (coefficient = -0.025, robust p-value = 0.286), local linear smoother fitted on each side of the discontinuity in dashed gray. Bins were created using mimicking variance evenly spaced method (Calonico et al., 2015). Each observation is a shooting victimization.

Figure 11: Sharp RD Estimates, by Bandwidth (100 - 2900m, in Increments of 200)

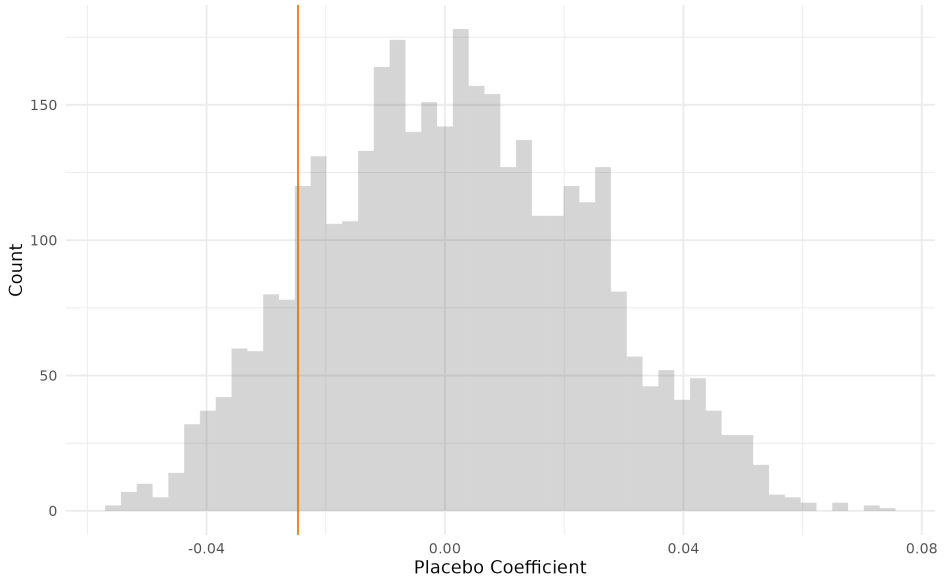


Note: 95% robust confidence intervals in gray. Estimates hold fixed a ratio of the h to b bandwidths to make estimates comparable to the baseline optimal bandwidth results. Appendix Figure A.3 presents equivalent estimates without this fixed ratio and qualitatively identical statistical significance.

VI.3 Additional Placebo Tests in Place

To provide an additional robustness check on the statistical inference of our baseline geographic RD estimate, we carry out a placebo test in geography (Figure 12) by estimating regression discontinuity models that follow our main specification at "pseudo-cutoffs" where no effect is expected. We find generally similar p-values for this test, assuaging concerns that our findings thus far are due to finite samples or other incorrect inferential assumptions. As the main identification strategy relies on ShotSpotter being discontinuously available on one side of a police district boundary versus another, we create placebo boundaries by shifting the existing boundary by a large number of randomly drawn constant values.[12] Under the placebo null distribution, our baseline estimate results in a p-value of 0.3, similar to that found in our baseline regression.

Figure 12: Placebo in Geography



Note: Baseline post-treatment RD estimate in orange.

[12]We constrain the placebo cutoff generation process to only keep draws that would allow enough of a sample size to estimate an effect (e.g., if a placebo cutoff would yield an in-bandwidth sample that has very few placebo control or placebo treatment observations, we do not count it towards the placebo null distribution). We generate 10,000 placebo cutoffs, then limit our placebo sample to those that would have at least 2,000 observations within a bandwidth equivalent to that of our baseline RD estimate, resulting in 3,558 valid placebo draws. Given that the choice of 2,000 is arbitrary, we assess the change in p-values to cutoffs other than 2,000, and find that between 1,000 and 5,000 observations, we would obtain p-values between 0.265 and 0.385.

VII Conclusions

Chicago, like many American cities, adopted ShotSpotter in part with the hope that the technology might help reduce gun violence.[13] Those hopes seem to have turned out to be in vain; an accumulating body of evidence suggests ShotSpotter does not seem to have any detectable impacts on criminal justice outcomes like shootings or arrests—though there is some evidence it leads to an increase in gun confiscations. Topper and Ferrazares (2024) suggest ShotSpotter may even have adverse spillover effects on criminal justice outcomes for other types of non-gun offenses, such as slower response times and fewer arrests, particularly for domestic violence cases.

The present study considers a different type of outcome from public health, which is also of great concern to society: whether a gunshot victim lives or dies. In the preliminary analysis we present here we use variation across geography (police districts) in the use of ShotSpotter technology; our design essentially compares whether there is a discontinuous jump in the shooting fatality rate at the district boundary. We find small (~1-2 percentage point) estimates during the period before ShotSpotter is implemented, which provides some support for the validity of the research design. During the period after ShotSpotter is in place, for different reasonable choices of bandwidth we see point estimates that imply 2 to 4 percentage point reductions in the shooting fatality rate, equal to around 14% to 25% of the baseline shooting fatality rate.

Our present research design does not allow for the level of statistical precision one might wish for. Even with such proportionately meaningful point estimates, the 95% confidence intervals allow for anything from very large reductions (-0.091) in the shooting fatality rate to modest increases (0.027) in fatality odds. We are currently working on additional analysis to incorporate variation over time within Chicago, with the hope that our multiple-regression-discontinuity design might improve the precision of our estimates.

But if we take our current point estimates at face value, what might they imply for public policy?

The cost side of the equation includes the $7-10 million the city of Chicago has averaged on the ShotSpotter contract (in 2024 dollars) and the possibility of slower response times and fewer arrests to non-gun 911 calls as suggested by Topper and Ferrazares (2024).

On the benefit side of the equation, with 1,966 outdoor shootings in Chicago in calendar year 2024 in those police districts that had ShotSpotter active, our baseline specification point estimate of a 2.5 percentage point reduction on shooting fatality rates would imply 49 fewer fatalities. If we were to value those lives using the $13 million figure the federal government uses for benefit-cost analysis of candidate safety regulation, our point estimates taken at face value would imply a benefit to society on the order of $637 million. Or another way to see it is to divide the financial cost of the policy by the number of lives saved, which would imply a cost per life saved of ~$171,000, which

[13] See for example Mayor's Press Office (2016).

compares very favorably to many, perhaps even most, other health and safety policies.[14]

The worthiness of this trade-off is ultimately a normative question left to policymakers. The main argument we seek to advance in this analysis is that any judgment about whether to retain or discard this new technology should be made with some consideration of all relevant outcomes beyond criminal justice ones, including public health outcomes.

[14]The ShotSpotter contract was extended in February 2024, covering a range from August 2018 to November 2024 and costing the city an annual average of $8.4M City of Chicago (2024a).

References

Banks, P. (2018, January). Re: First aid, Policy, training, and equipment. Policy Report, Civilian Office of Police Accountability, City of Chicago.

Black, S. E. (1999, May). Do Better Schools Matter? Parental Valuation of Elementary Education. *The Quarterly Journal of Economics 114*(2), 577–599.

Braga, A. A. and P. J. Cook (2018, July). The Association of Firearm Caliber With Likelihood of Death From Gunshot Injury in Criminal Assaults. *JAMA Network Open 1*(3), e180833.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014, November). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs: Robust Nonparametric Confidence Intervals. *Econometrica 82*(6), 2295–2326.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2015, October). Optimal Data-Driven Regression Discontinuity Plots. *Journal of the American Statistical Association 110*(512), 1753–1769.

Cattaneo, M. D., N. Idrobo, and R. Titiunik (2020). *A Practical Introduction to Regression Discontinuity Designs: Foundations*. Cambridge Elements: Quantitative and Computational Methods for Social Science. Cambridge: Cambridge University Press.

Cattaneo, M. D., M. Jansson, and X. Ma (2020, July). Simple Local Polynomial Density Estimators. *Journal of the American Statistical Association 115*(531), 1449–1455.

Cheng, A. (2024). Estimation of Regression Discontinuity and Kink Designs with Multiple Running Variables. *Working Paper*.

Cheng, Y. (2023, July). Houston's gunshot alert system isn't curbing violence but delays police response times, data shows.

Chicago Department of Transportation (2007, April). Street and Site Plan Design Standards. Design Standards, City of Chicago.

Chicago Police Department (2023). 2023 annual report. Technical report, Chicago Police Department, Chicago, IL.

City of Chicago (2024a, February). Agreement for shotspotter technology extension. https://news.wttw.com/sites/default/files/article/file-attachments/ShotSpotterFeb26Contract.pdf. Contract document between the City of Chicago and SoundThinking (formerly ShotSpotter, Inc.).

City of Chicago (2024b, October). Violence reduction - ShotSpotter alerts - historical. https://data.cityofchicago.org/Public-Safety/Violence-Reduction-Shotspotter-Alerts-Historical/3h7q-7mdb/about_data.

City of Chicago (2024c). Violence reduction - victims of homicides and non-fatal shootings. https://data.cityofchicago.org/Public-Safety/Violence-Reduction-Victims-of-Homicides-and-Non-Fa/gumc-mgzr/about_data. Dataset containing individual-level homicide and non-fatal shooting victimizations from 1991 to present.

Connealy, N.T., Piza, E.L., Arietti, R.A. et al (2024). Staggered deployment of gunshot detection technology in chicago, il: a matched quasi-experiment of gun violence outcomes. *Journal of Experimental Criminology*. Published March 23, 2024.

Cook, P. J. and A. Soliman (2024). Evaluation of Durham's ShotSpotter Installation: Results of a 12-Month Pilot Project. *SSRN Electronic Journal*.

Cunningham, R., M. Rosenburg, T. Corbin, C. Branas, S. Buggs, S. Haring, R. Jackson, A. Jain, J. Parsonnet, and B. Weston (2023, October). Using Science to Reduce Firearm Injuries and Deaths. *NAM Perspectives 10*(8).

Doucette, M. L., C. Green, J. Necci Dineen, D. Shapiro, and K. M. Raissian (2021, October). Impact of ShotSpotter Technology on Firearm Homicides and Arrests Among Large Metropolitan Counties: A Longitudinal Analysis, 1999–2016. *Journal of Urban Health 98*(5), 609–621.

Eckberg, D. L. (1995). Estimates of Early Twentieth-Century U.S. Homicide Rates: An Econometric Forecasting Approach. *Demography 32*(1), 1–16.

Ferguson, J. M. and D. Witzburg (2021, August). The Chicago Police Department's Use of Shotspotter Technology. Technical Report OIG FILE #21-0707, City of Chicago Office of Inspector General.

Goldenberg, A., D. Rattigan, M. Dalton, J. P. Gaughan, J. S. Thomson, K. Remick, C. Butts, and J. P. Hazelton (2019, December). Use of ShotSpotter detection technology decreases prehospital time for patients sustaining gunshot wounds. *Journal of Trauma and Acute Care Surgery 87*(6), 1253–1259.

Imbens, G. and K. Kalyanaraman (2012, July). Optimal Bandwidth Choice for the Regression Discontinuity Estimator. *The Review of Economic Studies 79*(3), 933–959.

Imbens, G. W. (2021, August). Statistical Significance, p -Values, and the Reporting of Uncertainty. *Journal of Economic Perspectives 35*(3), 157–174.

Koren, D. (2018, October). ShotSpotter pilot assessment.

Londoño-Vélez, J., C. Rodríguez, and F. Sánchez (2020, May). Upstream and Downstream Impacts of College Merit-Based Financial Aid for Low-Income Students: Ser Pilo Paga in Colombia. *American Economic Journal: Economic Policy 12*(2), 193–227.

Lopez, J. and T. Ballard (2024, May). Summer safety analysis: Chicago crime trends. Technical report, University of Chicago Crime Lab.

Ludwig, J. (2025). Unforgiving places: The unexpected origins of american gun violence. In *Unforgiving Places*. University of Chicago Press.

Ludwig, J. and J. Miller (2024, March). Jens Ludwig and Jacob Miller: We are seeing a lethal shift in America's gun violence crisis. *Chicago Tribune*.

Mares, D. and E. Blackburn (2021, June). Acoustic gunshot detection systems: A quasi-experimental evaluation in St. Louis, MO. *Journal of Experimental Criminology 17*(2), 193–215.

Mares, D. and A. Buettner (2021 (updated 2024)). The Impact of Shotspotter Deployment in Winston-Salem, NC. Technical report, Center for Crime Science and Violence Prevention.

Mayor's Press Office (2016, September). Public building commission approves expansion gun shot detection technology and POD cameras to fight crime in Englewood and Harrison districts.

McCrary, J. (2008, February). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics 142*(2), 698–714.

Nasser, A. A., C. Nederpelt, M. El Hechi, A. Mendoza, N. Saillant, P. Fagenholz, G. Velmahos, and H. M. Kaafarani (2020). Every minute counts: the impact of prehospital response time and scene time on mortality of penetrating trauma patients. *The American Journal of Surgery 220*(1), 240–244.

Piza, E. L., G. O. Mohler, J. G. Carter, D. N. Hatten, N. T. Connealy, R. Arietti, J. Cho, and E. Castillo (2024, January). The impact of gunshot detection technology on gun violence in Kansas City and Chicago: A multi-pronged evaluation. Technical Report 308357, National Criminal Justice Reference Service / U.S. Department of Justice, Office of Justice Programs.

Sampson, R. J. (2012). *Great American City: Chicago and the Enduring Neighborhood Effect*. The University of Chicago Press.

SoundThinking (2024). Gunshot Detection. https://www.soundthinking.com/law-enforcement/leading-gunshot-detection-system/.

Topper, M. and T. Ferrazares (2024). The Unintended Consequences of Policing Technology: Evidence from Shotspotter.

University of Chicago (2020, August). Chicago's Grid System. https://chicagostudies.uchicago.edu/grid.

Wellford, C. F. (2005). *Firearms and Violence: A Critical Review*. Washington, DC: National Academies Press.

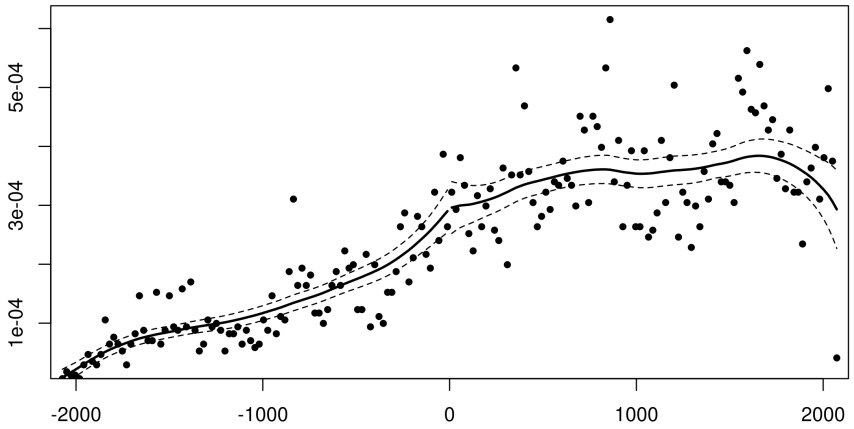
Wong, V. C., P. M. Steiner, and T. D. Cook (2013). Analyzing regression-discontinuity designs with multiple assignment variables: A comparative study of four estimation methods. *Journal of Educational and Behavioral Statistics 38*(2), 107–141.

Zimring, F. E. and G. Hawkins (1997, August). *Crime Is Not the Problem: Lethal Violence in America*. Oxford University Press.

A Appendix

A.I Appendix Figures

Figure A.1: McCrary (2008) Density Plot



Note: p-value = 0.983. Data includes only observations within two kilometers of the district boundary.

Figure A.2: Pre-Treatment Placebo Test, by Bandwidth (100 - 2900m, in Increments of 200)

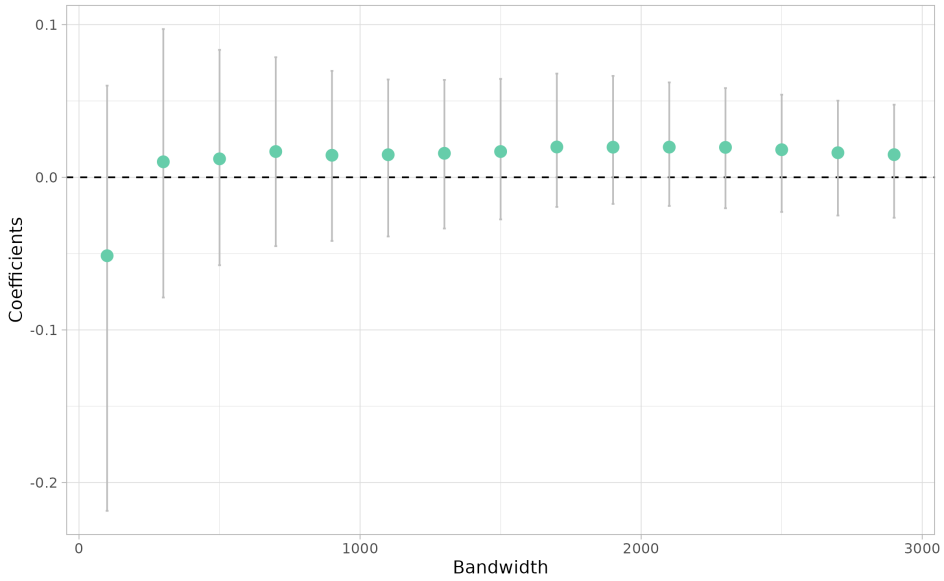
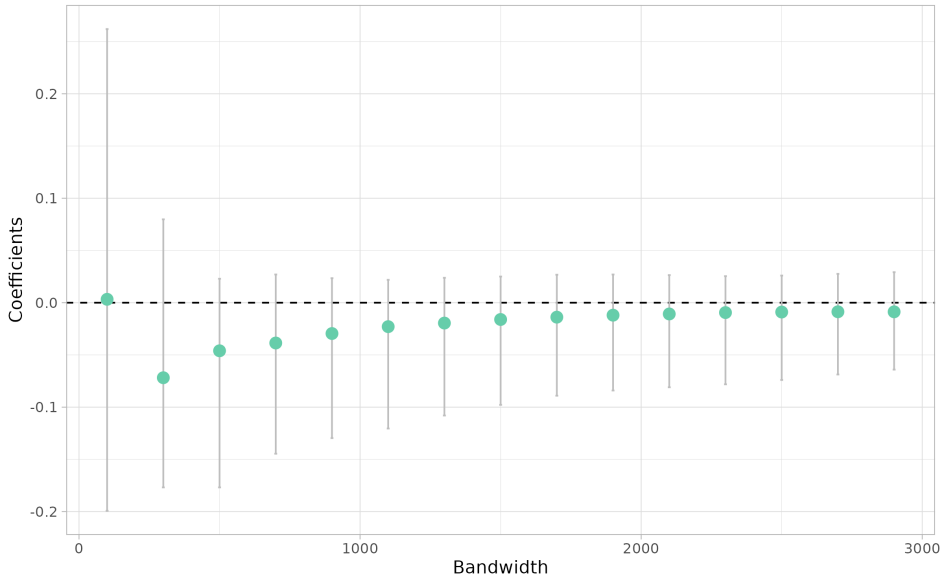


Figure A.3: Sharp RD Estimates, by 1:1 h/b Ratio Bandwidths (100 - 2900m, in Increments of 200)



A.II Appendix Tables

Table A1: Sharp Regression Discontinuity Estimates, by Polynomial and Kernel

	Outcome: Fatal Shooting		
In ShotSpotter Area	-0.018 [-0.083, 0.03] (0.353)	-0.036 [-0.109, 0.028] (0.246)	-0.034 [-0.106, 0.025] (0.228)
Optimal Bandwidth	883.73	1521.323	1369.125
Polynomial Order	1	2	2
Kernel	Uniform	Triangular	Uniform
Observations	3551	5568	5024

Note: Bandwidth selection robust confidence interval in brackets, p-value in parentheses.

Table A2: Continuity Tests of Victim/Incident Covariates, During-Treatment

Covariate	Control Mean	Treatment-Control Diff.	p-value	Observations
Age 0-19	0.217	-0.045	0.131	3957
Age 20-29	0.402	0.048	0.160	3957
Age 30-39	0.221	0.020	0.774	3957
Age 40-49	0.110	-0.019	0.461	3957
Age 50-59	0.034	0.010	0.405	3957
Age 60-69	0.016	-0.016	0.047	3957
Age 70-79	0.000	-0.000	0.250	3957
Age 80+	0.000	0.002	0.345	3957
Age (Missing)	0.015	0.010	0.126	4021
API/Other NH	0.003	0.008	0.068	3938
Black NH	0.659	0.082	0.019	3938
Hispanic/Latino	0.299	-0.083	0.013	3938
White NH	0.039	-0.007	0.542	3938
Race (Missing)	0.017	0.013	0.105	4021
Male	0.881	0.019	0.507	3969
Sex (Missing)	0.012	0.006	0.302	4021
Jan	0.063	-0.017	0.441	4021
Feb	0.046	-0.018	0.149	4021
Mar	0.041	0.023	0.143	4021
Apr	0.060	0.010	0.507	4021
May	0.108	0.012	0.334	4021
Jun	0.118	0.054	0.034	4021
Jul	0.126	-0.014	0.496	4021
Aug	0.107	0.035	0.183	4021
Sep	0.117	-0.007	0.915	4021
Oct	0.063	0.011	0.637	4021
Nov	0.074	-0.054	0.001	4021
Dec	0.076	-0.035	0.033	4021
Monday - Friday	0.598	-0.058	0.042	4021
CPD Watch #1	0.413	0.036	0.301	4021
CPD Watch #2	0.151	-0.029	0.127	4021
CPD Watch #3	0.437	-0.008	0.825	4021

Note: Each test carried out using the baseline specification outlined in equation (1), using a local linear regression, optimal bandwidths from the main RD, bandwidth selection robust inference, and a triangular kernel. Observation totals are comprised of treatment and control, within bandwidth.

A.III Appendix: Shooting to ShotSpotter (SST) Alert Matching

Due to the absence of unique identifiers between shooting victimization records and SST alert data, we perform spatial-temporal linking by identifying the SST alerts closest to shooting victimizations within defined time and distance parameters.

Spatial matches were conducted using within-distance thresholds between 200 and 450 meters to account for known data adjustments and SST sensor error. Thus, if an alert signaled that the gunfire originated within that radius of a shooting victimization, it would qualify as a matched pair. To determine these thresholds, we considered known data adjustment and sensor errors. According to SoundThinking, the owners of SST, the distance between the crime scene location and the location of where SST estimates the gunfire came from is typically within 25 meters (~82 feet) (SoundThinking, 2024). In addition, geospatial coordinates are deliberately jittered for both shooting and SST alert data by roughly the size of an average city block, with a new location randomly selected within the circumference of this circle (City of Chicago, 2024c,b). As a result, the locations of shootings and alerts may deviate by as much as two city blocks total (one block for each set of coordinates jittered), plus additional error from the alert data. Given that the average city block in Chicago is about 660 feet (~200 meters), we adjusted the matching distance parameters to range from within 200 meters to within 450 meters to account for this jittering and alert error. This theoretically captures the majority of relevant matches.

According to SoundThinking documentation (SoundThinking, 2024), the average time between a shooting and an alert being dispatched is approximately one minute. To balance capturing relevant matches while minimizing the risk of linking alerts to shootings at significantly different times, we adopted a less restrictive approach by exploring a time window of ± 5 to 20 minutes. These buffers assume the SST sensor-generated timestamps are more precise while accommodating potential inaccuracies in shooting timestamps. In practice, a narrow time difference within a reasonable geographic buffer is likely to provide more meaningful information than spatial distance alone, particularly given the previously noted coordinate jittering, and the fact that shootings in Chicago are highly geographically concentrated (Lopez and Ballard, 2024).

Although shooting incidents can generate multiple SST alerts, identifying victimizations that took up treatment only requires one reasonably matched alert. While the take up indicator would not change, we can select the alert that has the highest ranking using two methods: Time-Space Normalized, which scales time and distance values between 0 and 1, and Time-Space Normalized & Weighted, which does the aforementioned scaling exercise and assigns a specified weight to time and distance (i.e. 70% weight to the time difference and the remaining 30% to the spatial distance). The normalized ranking helps identify the single best match for each shooting incident and a deduplication step ensures no unique SST alert matches more than one shooting.

To evaluate how match rates vary with different constraints, we tested multiple combi-

nations of these parameters using the below criteria.

- The observation period is dynamic to begin on the SST district-level implementation dates as defined by Piza et al. (2024) and the cutoff for all districts is the day after SST stopped logging alerts following the Mayor's decision to end the contract. These dates are outlined in Appendix section A5.
- Although it is possible for an SST alert to estimate the origin of gunfire in a location of a non-SST district, the match rates only review alerts that matched to shooting cases that occurred in SST districts.
- Similarly, though it is possible for an alert to be matched to an indoor shooting incident, those cases are removed from the match rates analysis.

The overall match rates at the parameter-level (distance & time) are enumerated in the table below. The rates range from 51.3% to 56.9%, with an average rate of 54.6%. These results suggest that adjusting the match function parameters does not lead to a substantial variation in match quality across the treated districts. This outcome further implies that for every two outdoor or exterior shootings in an SST district, only one has a matched alert.

The average match rates across all parameters at the district-level are outlined in Table A4 below. These results indicate that there is variability on matching at the district-level, ranging from 47.6% in District 15 to 61% in District 5, a difference of 13.4 percentage points. An inspection of the average district-level match rates broken out by distance and time indicate that match rates within districts are also not largely impacted by toggling the distance and time window parameters as shown in the previous table.

An analysis using the parameters within 450 meter distance with a ± 15 minute time window provides descriptive evidence suggesting matching rate trends evolved over time, with varying levels and year-over-year directions observed at the district level (see Figure A.4 below). Generally, most districts have experienced a gradual increase in match rates since their activation, indicating an improvement in the likelihood of linking an alert to a shooting incident within a specified time and spatial range. However, total outdoor shooting incidents within districts do not consistently align with these match rate trends. For example, in District 11, outdoor shooting incidents decreased beginning in 2021, even as match rates showed a steady increase during the same period. This divergence highlights potential differences in trends between shooting incidents and SST alert accuracy or detection rates. We will use these findings to guide our Fuzzy RD work in the future.

Table A3: Match Rates by Distance and Time Thresholds

Shooting to Alert Distance (Meters)	Shooting to Alert Time (Mins)	Matched Shooting Cases	Total Shooting Cases	Match Rate
200	5	7246	14130	0.513
200	10	7455	14130	0.528
200	15	7554	14130	0.535
200	20	7602	14130	0.538
250	5	7401	14130	0.524
250	10	7614	14130	0.539
250	15	7717	14130	0.546
250	20	7770	14130	0.550
300	5	7488	14130	0.530
300	10	7705	14130	0.545
300	15	7816	14130	0.553
300	20	7869	14130	0.557
350	5	7545	14130	0.534
350	10	7766	14130	0.550
350	15	7882	14130	0.558
350	20	7938	14130	0.562
400	5	7584	14130	0.537
400	10	7808	14130	0.553
400	15	7925	14130	0.561
400	20	7984	14130	0.565
450	5	7629	14130	0.540
450	10	7855	14130	0.556
450	15	7977	14130	0.565
450	20	8042	14130	0.569

Table A4: Rollout Schedule and Match Rates by District

Rollout Date	District	Average Match Rate
2017-02-06	7	0.558
2017-02-06	11	0.489
2017-04-24	15	0.476
2017-07-14	9	0.550
2017-09-26	6	0.580
2017-10-30	10	0.576
2017-12-31	3	0.552
2018-01-31	4	0.600
2018-03-07	5	0.610
2018-04-11	8	0.501
2018-04-11	25	0.565
2018-05-16	2	0.552

Figure A.4: Annual Shooting Cases & SST-Alerts-to-Shooting-cases Match Rates



Note: Numerator and denominator records are filtered to SST Districts & Outdoor/Exterior cases only.

A.IV Appendix: ShotSpotter Implementation Dates

There are three main sources of information on ShotSpotter installation dates: Piza et al. (2024), Topper and Ferrazares (2024), and the ShotSpotter alert data from the City of Chicago data portal (City of Chicago, 2024b).

According to Piza et al. (2024, p. 25), there was a pilot phase of ShotSpotter in micro geographies within District 11 and 7 between 9/1/2012 and 2/5/2017. During this pilot phase, the system did not retain alert data. The remaining implementation across districts went through May 2018. It is unknown where this source of installation dates originated.

In Topper and Ferrazares (2024, p. 44), the rollout dates used span from 2017 – 2018. The authors credit a combination of the Chicago Police Department as well as Max Kapustin and Ashna Arora for assisting with details regarding rollout dates in Chicago.

The ShotSpotter alert data from the City of Chicago comes with alert dates and District labels. As such, it was possible to discern the earliest alert per district and to detect when alerts start to exceed an average weekly threshold of alerts. The latter analysis begins by filtering alert data to include only alerts occurring within ShotSpotter districts. To minimize driving up the average with potential false positives, alerts from New Year's weekends across all years are excluded, as alert counts typically spike significantly around this holiday. Weekly alert counts are then calculated at the district level for each week from the start of 2017 through present, with any null observations filled with a count of zero. Then, for each district, the average weekly counts are computed on an annual basis which serves as the average alert threshold. The data is then structured to indicate whether a week's alert counts meet or exceed the average alert threshold for each district-year. For districts that started using ShotSpotter in 2017, based on the Piza et al. (2024) dates, the 2018 average alert count serves as the threshold; for those implemented in 2018, the 2019 average is used. Finally, the earliest date at which the alert count meets or surpasses the threshold in each district is designated as the "rollout date" supplied in Table A5.

Table A5: Comparison of ShotSpotter Implementation Dates Across Sources

District	Piza et al. (2024) [1]	Topper & Ferrazares (2024) [2]	First Alert (City of Chicago) [3]	Weekly Alerts Exceed Average (City of Chicago) [4]	Days Between [1] - [2]	Days Between [1] - [3]	Days Between [1] - [4]
7	2017-02-06	2017-01-13	2017-01-13	2017-03-26	24	24	-48
11	2017-02-06	2017-03-01	2017-01-19	2017-04-16	-23	18	-69
15	2017-04-24	2017-05-13	2017-03-16	2017-05-14	-19	39	-20
9	2017-07-14	2017-06-01	2017-02-14	2017-07-09	43	150	5
6	2017-09-26	2017-09-24	2017-08-07	2018-02-11	2	50	-138
10	2017-10-30	2017-10-16	2017-03-19	2017-11-26	14	225	-27
3	2017-12-31	2018-01-04	2017-06-04	2018-02-04	-4	210	-35
4	2018-01-31	2018-02-01	2018-01-29	2018-01-28	-1	2	3
5	2018-03-07	2018-03-07	2018-03-13	2018-04-29	0	-6	-53
8	2018-04-11	2018-04-01	2017-01-21	2018-05-13	10	445	-32
25	2018-04-11	2018-04-24	2017-01-25	2018-04-29	-13	441	-18
2	2018-05-16	2018-05-16	2018-02-03	2018-05-20	0	102	-4

To identify the set of rollout dates to use for this analysis, we first analyzed the average absolute differences between available sources. The average discrepancy between Piza et al. and Topper & Ferrazares was 12.8 days, suggesting a high degree of consistency between these sources. In contrast, the alert data from the City of Chicago showed a larger average difference. For instance, comparing dates from Piza et al. with the initial alert date per district from the city data revealed an average difference of 142.7 days (over 4.5 months)—in all but one district, the first alert date was always earlier than the corresponding Piza et al. date. To assess the implications of this discrepancy, we calculated the number of alerts observed during this period, finding an additional 574 alerts (ranging from 1 to 196 in each district). Given that in 2019 there was an average of 1,586.25 citywide alerts per month (over 7,000 in 4.5 months), this relatively low alert volume recovered within the 4.5-month timeframe suggests ShotSpotter was not yet in full operation. Therefore, we concluded that the first alert date is not representative of a fully functional ShotSpotter deployment and excluded it as an implementation date.

Additionally, we compared the Piza et al. dates with the "alerts-exceed-average" method, which identified rollout dates based on when weekly alerts began to exceed the district's average. The average difference between these two approaches was 37.7 days (or 28.5 days when excluding an outlier in District 6, where the difference was 138 days). In this comparison, the alerts-exceed-average dates consistently followed the Piza et al. dates in all but two districts, suggesting we would be using a later set of rollout dates in general. Across these discrepancies, the alerts-exceed-average dates captured 827 alerts within roughly a month of delay—still lower than the monthly average of 1,586.25 alerts in 2019 but closer to an operational volume.

Ultimately, given the findings discussed, we opted to use the Topper & Ferrazares dates for the present study.

A.V Appendix: Multidimensional Regression Discontinuity (MRD) Designs

This appendix details the current state of the literature on how to estimate a regression discontinuity model when there are multiple running variables that determine treatment status. To set ideas, consider the setting of Londoño-Vélez et al. (2020), where college scholarships are awarded to prospective students who satisfy a grade point average requirement and whose family wealth is below a certain threshold. That is, $T = I(r_k \geq c_k \ \forall k)$, where each r_k represents a running variable (i.e., GPA and wealth), and each c_k represents their respective cutoffs for inclusion into treatment. This setting does not conform to standard regression discontinuity designs, as there are students for which $r_1 \geq c_1$ and $r_2 < c_2$, and vice-versa.

There are multiple approaches in these settings that still yield unbiased estimates for different treatment effect estimands, assuming all the relevant identifying assumptions are met (Wong et al., 2013). Two approaches that rely on standard regression discontinuity techniques are:

- Fuzzy regression discontinuity: From the perspective of condition 1, the students that do not meet condition 2 create non-compliance in treatment takeup. However, as long as condition 1 still induces a discontinuous jump in the probability of takeup, one can use instrumental variables to estimate a standard fuzzy regression discontinuity design. This method returns the LATE for those at the cutoff for condition 1 who were induced to takeup by crossing the condition 1 threshold. One can similarly apply the same idea to condition 2 and estimate the other LATE.
- Separate estimation: As in Londoño-Vélez et al. (2020), one can also subset the data to the sample for which one condition is met, and estimate the sharp regression discontinuity estimate for the other condition. In their case, they subset the data to all students that satisfy the family wealth requirement, and proceed to estimate the usual RD using GPA as a running variable, and vice-versa.

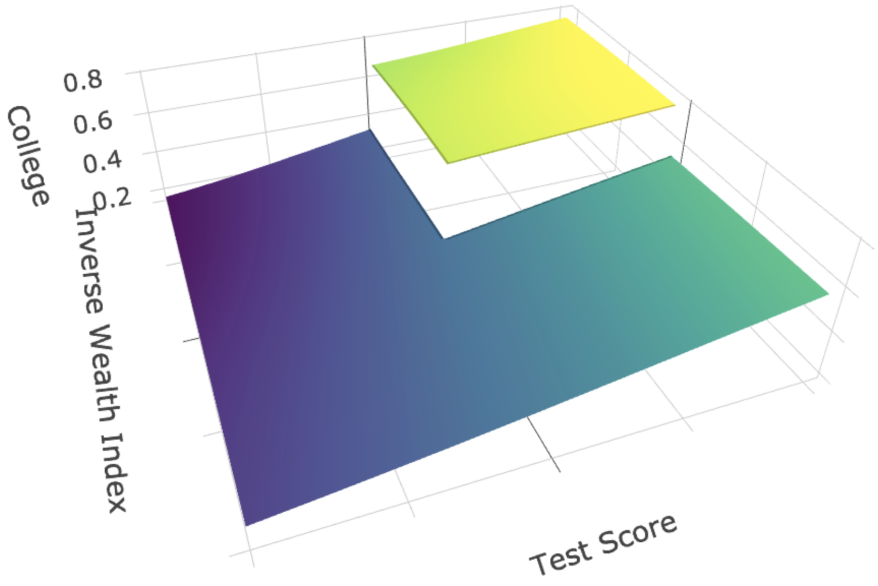
There are two pitfalls to these approaches—conceptually, one is simplifying a multidimensional problem into two single dimensional ones. The resulting estimates are then only relevant for the specific margin they're estimated at as opposed to the overall "frontier" that determines treatment. That is, one estimates the effect of satisfying the GPA cutoff or the wealth cutoff, but not both simultaneously, which is the actual policy at hand. The second pitfall is a statistical power one, as the data is being split into subsets to generate each estimate, effectively leaving useful variation in the data on the table.

While we will still create these estimates as a robustness check, we instead propose to leverage newer econometric methods that allow us to use all the available data to estimate treatment effects local to the entire frontier, with an efficiency gain over existing methods. Our approach is to leverage the insights from Cheng (2024) to estimate treatment effects from an MRD using thin-plate splines.

The general intuition from Cheng (2024) is that just as one can represent a standard RD in a two-dimensional graph with the running variable on the x-axis and the outcome as a function of x on the y-axis, one can represent an MRD as a three-dimensional graph where the x- and y-axes represent two running variables, and the z-axis represents the outcome as a function of x and y. Then, the quadrant defined by the cutoffs being met represents the treatment region (Figure A.5), and one can plot a surface that shows the conditional expectation of the outcome as both running variables vary.

The remaining hurdle is to then devise a method to estimate this conditional expectation. In the standard RD setting, the typical approach is to estimate two local linear regressions, one on each side of the cutoff. The difference between those two lines at the cutoff point yields the causal estimand of interest. Cheng (2024) proposal is to extend that logic by estimating two thin-plate splines, one for the treated quadrant, and one for the control quadrant, and take their vertical difference at any point along the treatment frontier to obtain a conditional average treatment effect (CATE) at that specific point, or to average over the entire frontier to obtain a frontier average treatment effect.

Figure A.5: Graphical Representation of an MRD



Note: from Cheng (2024).

A useful implication from this approach is that one can verify this method can yield CATEs that are consistent with the cutoff specific approaches listed above—one estimates the overall model (e.g., for GPA and wealth combined), then averages over the CATEs for a specific frontier (e.g., for GPA only), and compares that estimate to the single dimensional RD for GPA as a robustness check. For more details on the thin-plate spline estimation procedure itself, see Cheng (2024).